

1/25° Atlantic Ocean Simulation Using HYCOM

Alan J. Wallcraft¹, Eric P. Chassignet², Harley E. Hurlburt¹, and Tamara L. Townsend¹

¹ Naval Research Laboratory, Stennis Space Center, MS 39529

² Rosenstiel School of Marine and Atmospheric Science, U. Miami, Miami, FL 33149

Introduction

Traditional ocean models use a single coordinate type to represent the vertical, but no single approach is optimal for the global ocean (Chassignet et al., 2000; Willebrand et al., 2001). Isopycnal (density tracking) layers are best in the deep stratified ocean, z-levels (constant fixed depths) provide high vertical resolution in the mixed layer, and terrain-following levels are often the best choice in coastal regions. The HYbrid Coordinate Ocean Model (HYCOM) (Bleck, 2002) combines all three approaches by dynamically choosing the optimal distribution at every time step via the layered continuity equation. This has led to HYCOM being chosen for the next generation of ocean prediction systems both by NAVOCEANO and by NOAA.

A prototype near real-time 1/12° Atlantic HYCOM prediction system is already running in NAVOCEANO in demonstration mode with assimilation of analysis fields derived from satellite data, sea surface temperature and altimetric sea surface height (Chassignet et al., 2005a). Results are at <http://hycom.rsmas.miami.edu>. This system has demonstrated the capability to map the “ocean weather”, including the meandering of ocean fronts and currents (e.g. the Gulf Stream) and oceanic eddies. Comparison of five ocean prediction systems with ocean color imagery in the Gulf of Mexico proved very effective in discriminating between the performance of these systems in mapping the Loop Current and associated eddies (Chassignet et al., 2005b). These five systems, one using HYCOM, has horizontal resolutions of 4, 6, 8, 8, and 17 km in the Gulf of Mexico. The 1/12° Atlantic HYCOM system with 8 km resolution was outperformed only by the system with double the horizontal resolution, comparable to the 1/25° Atlantic HYCOM described here. The prediction system with 4 km resolution uses a computationally less expensive ocean model, but it excludes shallow water. HYCOM includes shallow water and is particularly well designed for transition between deep and shallow water, historically a challenging problem for ocean models.

HYCOM’s basic parallelization strategy is 2-D domain decomposition, *i.e.*, the region is divided up into smaller sub-domains, or tiles, and each processor “owns” one tile. Tiles that are entirely over land are discarded (figure 1). A halo is added around each tile to

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE DEC 2005		2. REPORT TYPE		3. DATES COVERED 00-00-2005 to 00-00-2005	
4. TITLE AND SUBTITLE 1/25degree Atlantic Ocean Simulation Using HYCOM				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory,Stennis Space Center,MS,39529				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES 9th HYCOM Consortium Meeting, Dec 6-8, 2005, RSMAS, Miami, FL					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 8	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			




Figure 1: HYCOM 1/25° Atlantic model domain decomposition for 507 MPI tasks

allow communication operations (*e.g.*, updating the halo) to be completely separated from computational kernels, greatly increasing the maintainability and expandability of the code base. Rather than the conventional 1 or 2 element wide halo, HYCOM has a 6 element wide halo which is “consumed” over several operations to reduce halo communication overhead. The communication mechanism implementing domain decomposition can be either MPI or Cray’s SHMEM library.

For basin-scale applications it is important to avoid calculations over land. MICOM was a pioneer in optimizations that avoid land (Bleck et al., 1995). It fully “shrink wrapped” calculations on each tile and discarded tiles that were completely over land. HYCOM goes farther than MICOM, and most other structured grid ocean models, in land avoidance by allowing more than one neighboring tile to the north and south. Figure 1 illustrates an Atlantic tiling with equal sized tiles that a) allows rows to be offset from each other if this gives fewer tiles over the ocean and b) allows two tiles to be merged into one larger tile if less than 50% of their combined area is ocean. The North African region of figure 1 illustrates both these optimizations. A conventional 4 neighbor equal sized tiling would use 560 MPI tasks out of the 990 original tiles (30 by 33), but we only need 507 MPI tasks. More memory is required on some tiles and the communication

overhead is slightly increased, but the 10% saving in MPI tasks is a significant optimization given the computer requirements of this application.

Scalability Study

HYCOM routinely runs on 500 to 750 processors, but under Phase I of HPCMP's FY04 Capability Applications Projects (CAP) program we explored scalability to 2,000 processors both Globally at $1/12^\circ$ (~7 km mid-latitude) resolution (array size 4500x3298x26) and in the Atlantic at $1/25^\circ$ (array size 3357x3111x28). The first machine to become available was an IBM P655 (Power 4+) at NAVO with 2,944 processors (350 eight processor nodes). Since Global HYCOM is included in the HPCMP TI-05 benchmark suite, we started with the TI-05 case, but because a typical run is at least 30x longer than the benchmark we ignored the startup time before the first model time step. These initial runs identified three areas of suboptimal scalability to large processor counts: a) halo exchanges, b) global sums, and c) I/O. Of these, halo exchanges are not easy to optimize further (and the poor scalability may partly be due to aliased load imbalance, i.e. improving load balance may be more important than speeding up halo communications). For reasons unrelated to scalability, the latest version of HYCOM uses many fewer global sums than the benchmark code. Simply updating to the new version improved 504 to 1006 scalability (ratio of wall times) from 1.37x faster to 1.6x faster, and the speedup was almost entirely due to a reduction in the number of global sums. The final scalability issue is I/O.

HYCOM does I/O one 2-D array at a time. The model uses REAL*8 internally, but all I/O is REAL*4 and must be "big-endian" (so that the files are machine-independent). At present all I/O is performed by the first task, which gathers the array and then writes it out. Single-task I/O can be a bottleneck, but each I/O request from the Global domain is 56.6 MB which helps improve performance. HYCOM was doing its array gathering onto the I/O task on the original REAL*8 arrays, but since the I/O is REAL*4 the amount of data transferred was halved by switching to REAL*4 earlier in the process. This reduced I/O time by about a third.

On the IBM P655, our final scalability timing was:

MPI TASKS	NODES	WALL-TIME	SPEED-UP
504	63	1515.1	
1006	126	946.9	1.60x 504
2040	255	587.2	1.61x1006

The total I/O time was between 88 to 96 seconds and without I/O the 1006 to 2040 speedup would be 1.74x. The 2040 task case is spending 15% of its time doing I/O.

However, the test case is only half a model day and this amount of I/O would be more typical for a full model day.

The second machine to become available to this project was the Linux Networx Evolocity II commodity cluster system at ARL. It has 1024 dual-processor (3.4 GHz Intel Xeon EM64T) compute nodes, each with 4 GB of memory. On this system our global scalability timing was:

MPI TASKS	NODES	WALL-TIME	SPEED-UP
504	252	1867.0	
1006	503	1209.2	1.54x 504
2040	1020	772.1	1.57x1006

The total I/O time is between 336 to 284 seconds and without I/O the 1006 to 2040 speedup would be 1.84x. The 2040 task case is spending 35% of its time doing I/O.

MPI-2 I/O is an obvious alternative to HYCOM's single-task I/O, but the HYCOM arrays contain "holes" over land and we want these regions to be filled by a "data_void" value. MPI-2 I/O can handle I/O with gaps, but it can't fill the gaps with a data_void. So the best alternative appears to be to have one task in each row of the 2-D task decomposition do the I/O for all tasks in its row (after filling in any data voids). The I/O can then occur (on some machines) in parallel, but each I/O request will be smaller. For example, 29 and 33 tasks would be used for I/O out of 504 and 1006 tasks respectively - so each I/O request is now about 1.7-1.9 MB. These are probably still big enough that MPI-2 need not aggregate I/O requests on to fewer tasks. We setup a HYCOM I/O benchmark to test various I/O strategies (it is possible that different MPI-2 based approaches will be best on different machine architectures). We found some improvement in read performance using MPI-2 I/O, but little improvement in write performance (and most HYCOM I/O is writes). This was the case on both the IBM P655 and the Linux Networx Evolocity II commodity cluster.

Due to time constraints, we did not run the 1/25° Atlantic case on the IBM P655, but on the Xeon cluster we get (wall time for one model day):

MPI TASKS	NODES	WALL -TIME	SPEED-UP
507	254	2620.1	
1006	503	1588.1	1.65x 507
2017	1009	1027.6	1.55x1006
387	194	3250.6	
757	379	1959.3	1.66x 387
1533	767	1198.4	1.63x 757

The 1/25° Atlantic domain is 3357x3111x28, about 76% of the 1/12° Global domain (4500x3298x26), so the I/O time is slightly reduced (about 260 seconds) and is now once per one model day as is typical for actual runs. It is still 25% of the total time on 2017 processors. Allowing for checkpoint file I/O and startup, a typical model-month job would take 13.8 wall hours on 1006 cpus or 167K cpu hours per model year.

1/25° Atlantic Ocean Simulation

Under CAP at ARL, we have approval to run the 1/25° Atlantic Ocean simulation for five model years on 1006 processors. This is the first basin-scale demonstration of our target resolution for the global ocean, made possible several years ahead of schedule by CAP.

The simulation started from an existing 1/12° Atlantic run that has spun-up nine years starting from climatology. Figure 2a shows the sea surface height from the initial 1/12° state, and figure 2b shows the sea surface height from the 1/25° simulation after 13 model months. Both plots are in “logical” array space, i.e. each grid point maps to a square cell on the plot. The plot projections differ above 47°N because the 1/12° grid is Mercator everywhere but the 1/25° grid is a subset of the Global grid which has an Arctic dipole patch above 47°N.

As of this writing, the simulation has not run long enough to be sufficiently spun-up for substantial analysis. Based on experience with less expensive models (Hurlburt and Hogan, 2000; Shriver et al., 2005), the most striking change in sea surface height (SSH) from 1/12° to 1/25° resolution should occur in the Gulf Stream region extending between the separation of the stream from the coast near Cape Hatteras, North Carolina, and south of the Grand Banks (SE of Newfoundland) near 45°W. In particular, the model should develop a more robust nonlinear recirculation gyre, marked by high SSH, along the south side of the Gulf Stream with a tongue of low SSH separating it from the remainder of the high SSH in the subtropical gyre to the south. This changes the large scale circulation of the subtropical gyre giving it a distinct C-shape. Figure 2a (the initial state from the 1/12° model) shows little evidence of the C-shape, but it is beginning to emerge in the 1/25° simulation (figure 2b) marked by a series of anticyclonic eddies (high SSH) extending to 47°W between 35° and 40°N. This is separated from the remainder of the subtropical gyre to the south by a series of cyclonic eddies (low SSH) between 32° and 36°N. At 1/12° resolution getting a realistic pathway for the Gulf Stream between Cape Hatteras and the Grand Banks, such as is seen in figure 2a, is a challenge for HYCOM and other models, one that requires careful overall experimental design and substantial

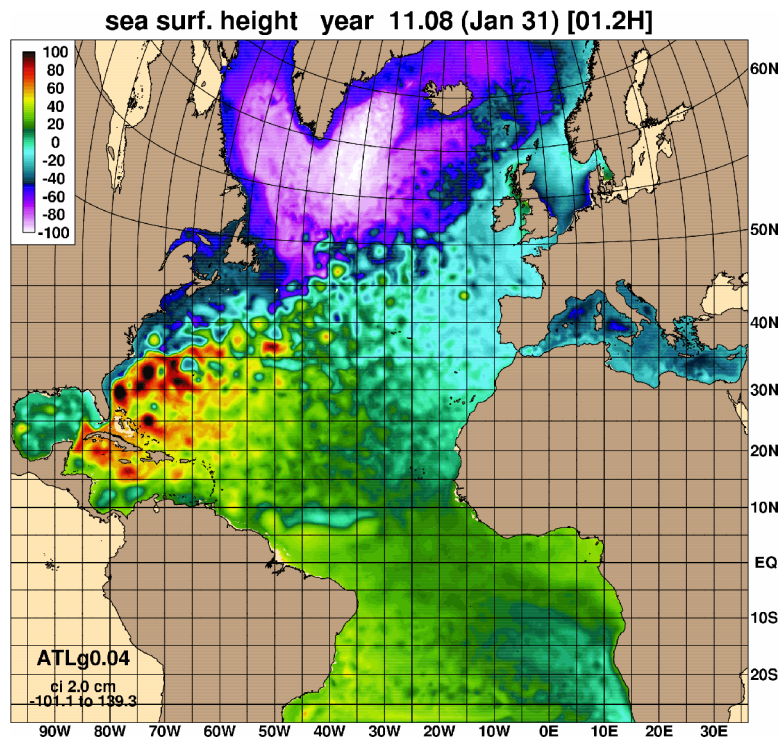



Figure 2: Sea Surface Height (cm) from (a) the initial $1/12^\circ$ state, and (b) the $1/25^\circ$ simulation after 13 model months.

tuning of the model parameters. At $1/25^\circ$ resolution we anticipate a more robust result, consistent with those in Hurlburt and Hogan (2000) and Shriver et al. (2005).

An increase in the amount and strength of the eddy activity throughout much of the model domain is another change that is clearly evident with the resolution increase in figure 2. The eddying and meandering of currents are associated with flow instabilities. A type of flow instability known as baroclinic instability is very efficient in pumping energy from the upper ocean to the abyssal ocean. In turn, abyssal flows can have a large effect in steering the pathway of upper ocean currents, making the simulation of current pathways more accurate. This has been demonstrated for the Kuroshio south and east of Japan (Hurlburt et al., 1996; Hurlburt and Metzger, 1998), the Tasman Front between Australia and New Zealand (Tilburg et al., 2001) and the Japan East Sea (Hogan and Hurlburt, 2000 and 2005), for example.

Other advantages of the resolution increase are better representation of straits and narrows (choke points) and shallow coastal regions. Resolution of $1/25^\circ$ (3-4 km at mid-latitudes) is sufficient for useful results in coastal regions worldwide. Furthermore, this resolution would allow direct nesting of coastal models with ~ 1 km resolution in Global HYCOM, largely eliminating the need for nested regional models. Of particular significance for nested coastal models is the demonstration by Shriver et al. (2005) that a model with this resolution can map small eddies (25-75 km in diameter) with about 70% reliability by assimilation of satellite altimeter SSH data. At $1/12^\circ$ resolution HYCOM would be limited to mapping eddies greater than 50 km in diameter.

Acknowledgments: This work was supported in part by a grant of computer time from the DOD High Performance Computing Modernization Program at NAVO MSRC and ARL MSRC. It was sponsored by the National Ocean Partnership Program (NOPP), and the Office of Naval Research (ONR) through the following projects: NOPP U.S. GODAE: Global Ocean Prediction with the HYbrid Coordinate Ocean Model (HYCOM), 6.1 NRL (ONR): Global Remote Littoral Forcing via Deep Water Pathways.

References

- Bleck, R., 2002. An oceanic general circulation model framed in hybrid isopycnic-cartesian coordinates. *Ocean Modeling*, 4, 55-88.
- Bleck, R., S. Dean, M. O'Keefe and A. Sawdey, 1995. A Comparison of Data-Parallel and Message-Passing Versions of the Miami Isopycnic Coordinate Ocean Model. *Parallel Computing*, 21, 1695-1720.
- Chassignet E.P., H. Arango, D. Dietrich, T. Ezer, M. Ghil, D.B. Haidvogel, C.-C. Ma, A. Mehra, A.M. Paiva, Z. Zirkles, 2000. DAMEE-NAB: The base experiments. *Dyn. Atmos. Oceans*, 32, 155-183.

Chassignet, E.P., H.E. Hurlburt, O.M. Smedstad, G.R. Halliwell, P.J. Hogan, A.J. Wallcraft, R. Baraille, and R. Bleck, 2005a. The HYCOM (HYbrid Coordinate Ocean Model) data assimilative system. *J. Mar. Sys.*, in press.

Chassignet E.P., H.E. Hurlburt, O.M. Smedstad, C.N. Barron, D.S. Ko, R.C. Rhodes, J.F. Shriver, A.J. Wallcraft and Robert A. Arnone, 2005b. Assessment of Data Assimilative Ocean Models in the Gulf of Mexico using Ocean Color. In AGU monograph “*New Developments in the Circulation of the Gulf of Mexico*”, submitted.

Hogan, P.J. and H.E. Hurlburt, 2000. Impact of upper ocean - coupling and isopycnal outcropping in Japan/East Sea models with $1/8^\circ$ to $1/64^\circ$ resolution. *J. Phys. Oceanogr.*, 30, 2535-2561.

Hogan, P.J. and H.E. Hurlburt, 2005. Sensitivity of simulated to surface wind forcing in the Japan/East Sea. *Deep Sea Res.* (in press)

Hurlburt, H.E. and P.J. Hogan, 2000. Impact of $1/8^\circ$ to $1/64^\circ$ resolution on Gulf Stream model-data comparisons in basin-scale subtropical Atlantic Ocean models. *Dyn. Atmos. Ocean.*, 32, 283-329.

Hurlburt, H.E. and E.J. Metzger, 1998. Bifurcation of the Kuroshio Extension at the Shatsky Rise. *J. Geophys. Res.*, 103, 7549-7566.

Hurlburt, H.E., A.J. Wallcraft, W.J. Schmitz Jr., P.J. Hogan, and E.J. Metzger, 1996. Dynamics of the Kuroshio/ Oyashio current system using eddy-resolving models of the North Pacific Ocean. *J. Geophys. Res.*, 101, 941-976.

Shriver, J.F., H.E. Hurlburt, O.M. Smedstad, A.J. Wallcraft and R.C. Rhodes, 2005. $1/32^\circ$ real-time global ocean prediction and value-added over $1/16^\circ$ resolution. *J. Mar. Sys.* (accepted)

Tilburg, C.E., H.E. Hurlburt, J.J. O'Brien, and J.F. Shriver, 2001. Dynamics of the East Australian Current system: the Tasman Front, the East Auckland Current and the East Cape Current. *J. Phys. Oceanogr.*, 31, 2917-2943,

Willebrand, J., B. Barnier, C. Boening, C. Dieterich, P. Herrmann, P.D. Killworth, C. LeProvost, Y. Jia, J.-M. Molines, and A. New, 2001. Circulation characteristics in three eddy-permitting models of the North Atlantic. *Prog. Oceanogr.*, 48, 123-161.